## 32.1 Architectures and Circuits for Software-Defined Radios: Scaling and Scalability for Low Cost and Low Energy

L. Van der Perre[1], B. Bougard[1], J. Craninckx[1], W. Dehaene[2], L. Hollevoet[1], M. Jayapala[1], P. Marchal[1], M. Miranda[1], P. Raghavan[1], T. Schuster[1], P. Wambacq[1], F. Catthoor[1], P. Vanbekbergen[1]

[1]IMEC, Leuven, Belgium, [2]K.U.Leuven, Leuven, Belgium

Wireless terminals supporting an increasingly large variety of modes (Cellular, WiFi, WiMax, WPANs) are subject to a cost increase that can be overcome by Software-Defined Radios (SDRs). Energy efficiency, however, is the main obstacle to successfully deploying reconfigurable radios. To address this, the focus here is on the design of energy scalable [1] architectures and circuits for SDRs. A system-level quality/energy manager applies cross-layer optimization. It exploits the scalability and the dynamics in application requirements and propagation environment to realize low power operation (Fig. 32.1.1).

The design of the reconfigurable RF front-end faces widely varying requirements. Carrier frequencies should range from a few 100MHz up to 6GHz and higher, channel bandwidths vary by 2 orders of magnitude, and linearity requirements are relaxed for constant-envelope systems but are very tough for OFDM. Moreover, in order to enable cross-layer optimization and control, energy scalability must be imposed [6] (see Fig. 32.1.1). This has serious consequences for the design of all building blocks in the front-end.

For example, the ADC must work at state-of-the-art performance levels and energy constraints for sampling rates ranging from kHz to 10's-of-MHz. The design [3] achieves full performance/energy scalability, with power consumption proportional to the sampling frequency from 50MS/s down to 0.

Also at the transmitter, a large gain in energy efficiency is realized when the linearity and output power is tuned to the instantaneous requirements instead of to the worst-case. Figure 32.1.2 shows a 3-stage PA driver (coupled with an external PA), which is made reconfigurable by changing load resistance, degeneration levels and bias currents [2]. The driver gain can be controlled over the range of -20 to +16dB, IIP3 ranges from +2 to +14dBm, while its current consumption scales with performance from 7 to 61mA. Transmit optimization, capitalizing on this flexibility, enables significant energy scalability. Considering a path-loss of 80dB for instance (see Fig. 32.1.3), the transmitter energy consumption is 25nJ/b when configured to transmit at 36Mb/s (16QAM, code rate 3/4), while it scales down to 16.6nJ/b (33% improvement) when configured for 24Mb/s transmission (16QAM, code rate 1/2).

The digital baseband platform achieves both functional flexibility and energy efficiency through opportunistic partitioning (Fig. 32.1.4). The sub-functions are divided according to their nature (control/DSP) and flexibility/energy efficiency requirements. The heterogeneous multi-processor SoC includes a Digital Baseband Front-End (DBFE) to implement packet detection. The remainder of the platform, which is heavily duty-cycled, consists of two baseband processors for data processing and an ARM926 for event-based control flow, inter-core data transfer, and MAC.

Analysis of typical wireless modem use cases shows that the power consumption in idle time often dominates battery life-time. Therefore, an ultra-low power DBFE is designed to perform all listening functionality, featuring dedicated flexibility (Fig. 32.1.4). The receiver DBFE carries out AGC, signal decimation, detection and coarse synchronization for up to 3 parallel concurrent streams. Buffering is provided so as not to loose data while proceeding to detection/synchronization. The architecture supports hierarchical activation. The AGC circuitry is always on and when the receiver detects power the decimation filter, circular buffer and detection/synchronization processor are activated. The decimation filters (40MS/s) cover current standards within a power budget of 0.5mW.

For the very demanding task of synchronization processing, a dedicated processor is designed. It exploits data parallelism yielding a dual-pipeline scalar/16-slot vector architecture. Since the effective number of instructions is reduced and the data-path is fully utilized by the SIMD (single instruction multiple data) approach, considerable energy is saved. A distributed loop buffer further reduces the instruction fetch overhead. The energy efficiency has been evaluated to be 0.1mW/MHz after physical synthesis using Synopsys PrimePower™, for detection of 802.11a/g.

The circular buffer is 8Kb (detection requires a window of 512 samples, at 16b/sample). A fine-grain leakage-controlled SRAM architecture is designed [4], using dual supply in the SRAM matrix. First there is the nominal supply for the periphery and the memory cells that are read or written. The second supply is applied to all the other cells in the matrix that are not activated and just have to maintain their state. Due to drain induced barrier lowering (DIBL), the leakage of the transistors is dependent on the supply voltage. A lower supply voltage reduces the matrix leakage (Fig. 32.1.5). In the design, the finest possible granularity of matrix activation is used; only the 16 bits that are really activated are switched to the higher supply. The voltage of the lower supply is determined by the minimum noise margin that is required for non-active cells, usually between 50 and 100mV. A specific circuit automatically regulates the lower supply as a function of this noise margin and the process parameters. Notably, the supply voltage for the periphery and the active cells in the matrix does not need to be the nominal supply. If maximum speed is not required, it can also be lowered.

Active power consumption is limited by variability. Instead of using the worst-case corner for timing closure, small-grain adaptive circuitry is adopted. The dynamic power of the SRAM depends strongly on its required access latency, normally imposed by the worst-case system load of the SDR, while the average requirements are often more relaxed. The SRAM design exploits this dynamic access latency versus energy trade-off at run-time with a control knob to decide the working point. It is realized by replacing the buffers inside the pre-decoder, word-line driver and post-decoder with switcheable (tapered) ones, being a combination of a slow and energy-efficient driver with a fast, but higher power one [5].

A switcheable decoder for a 8Kb memory has been designed and laid out in a 130nm CMOS technology. It achieves a 30% energy savings in exchange for 28% longer access latency. The configurable memory can be exposed to the system-level quality/energy manager.

The SDR described can effectively achieve low power operation if complemented by appropriate run-time adaptation. Wireless systems are facing very dynamic radio propagation conditions and QoS requirements (Fig. 32.1.1). Traditional designs are still mostly tuned for the worst-case. However, by carefully monitoring and adapting to the dynamics of the system, building on energy scalable SDR solutions and applying cross-layer optimization, much energy can be saved. For example, the proposed transmitter presents an energy-scalability range up to 30%. By continuously adapting to QoS requirements, this is translated into an average system-level energy efficiency improvement of up to 40%.

*References:*

[1] A. Sinha , A. Wang, and A. P. Chandrakasan, "Energy Scalable System Design," *IEEE Trans. VLSI Systems*, vol. 10, pp. 135-145, April, 2002.

[2] B. Debaillie, et al.," Energy-Scalable OFDM Transmitter Design and Control," *DAC*, pp. 536-541, July, 2006.

[3] J. Craninckx, et al., "A 65fJ/Conversion-Step, 0-50MS/s 0-0.7mW 9bit Charge-Sharing SAR ADC in 90nm Digital CMOS," *ISSCC Dig. Tech. Papers*, pp. 246-247, Feb., 2007.

[4] Peter Geens, et al., "A Small Granular Controlled Leakage Reduction System for SRAMs," *J. Solid-State Electronics*, vol. 49, pp. 1776-1782, Nov., 2005.

[5] H.Wang, et al., "Variable Tapered Pareto Buffer Design and Implementation Allowing Run-Time Configuration for Low Power Embedded SRAMs," *IEEE Trans. VLSI Systems*, vol. 13, pp. 1127-1135, Oct., 2005.

[6] J. Craninckx, et al., "A Fully Reconfigurable Software-Defined Radio Transceiver in 0.13µm CMOS," *ISSCC Dig. Tech. Papers*, pp. 346-347, Feb., 2007.

 

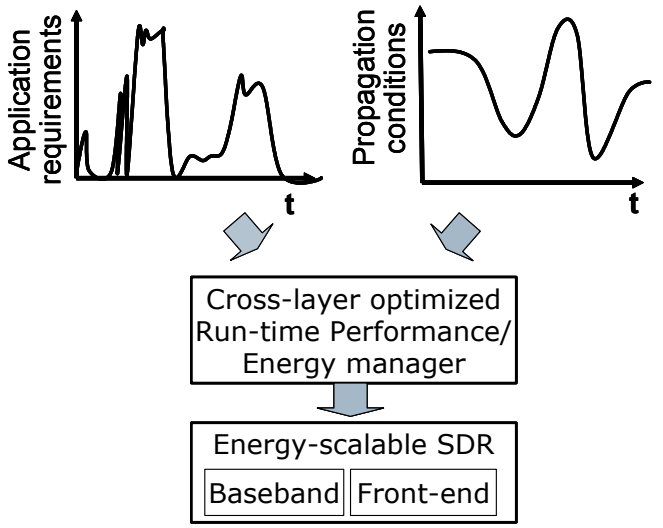


Figure 32.1.1: Energy scalable SDRs achieve low power operation through cross-layer joint QoS and energy management.

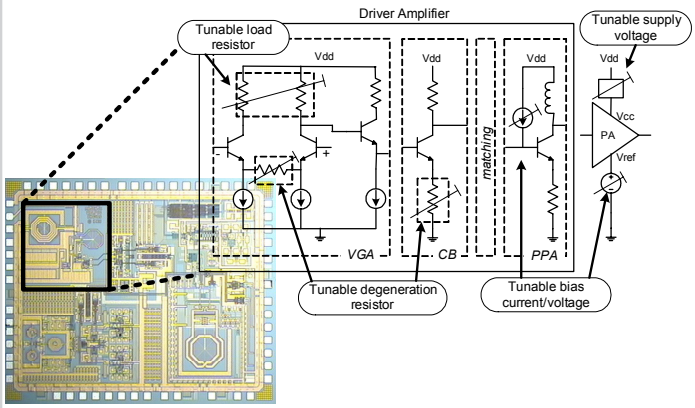


Figure 32.1.2: Energy scalable OFDM transmitter chain.

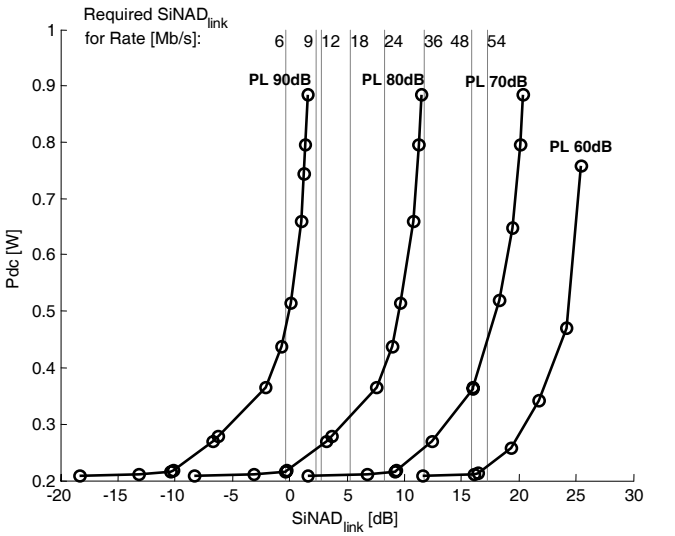


Figure 32.1.3: Power consumption of the transmitter front-end (Pdc) versus SiNAD$_{link}$ trade -off for various path-loss.

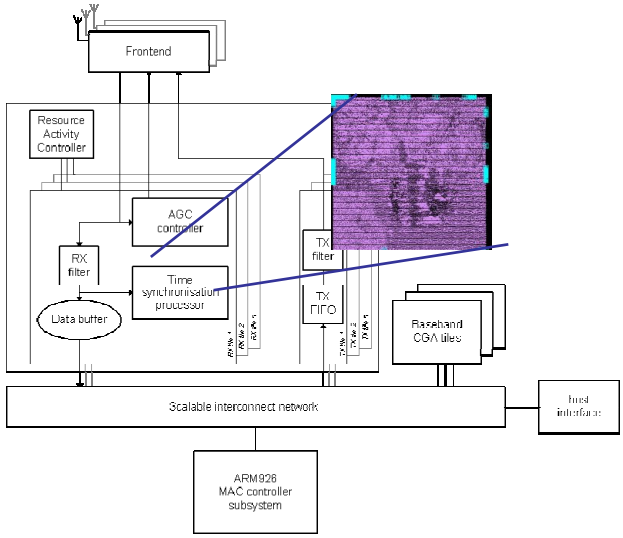


Figure 32.1.4: Heterogeneous MP SoC for SDR baseband platform including the Digital Baseband Front-End.

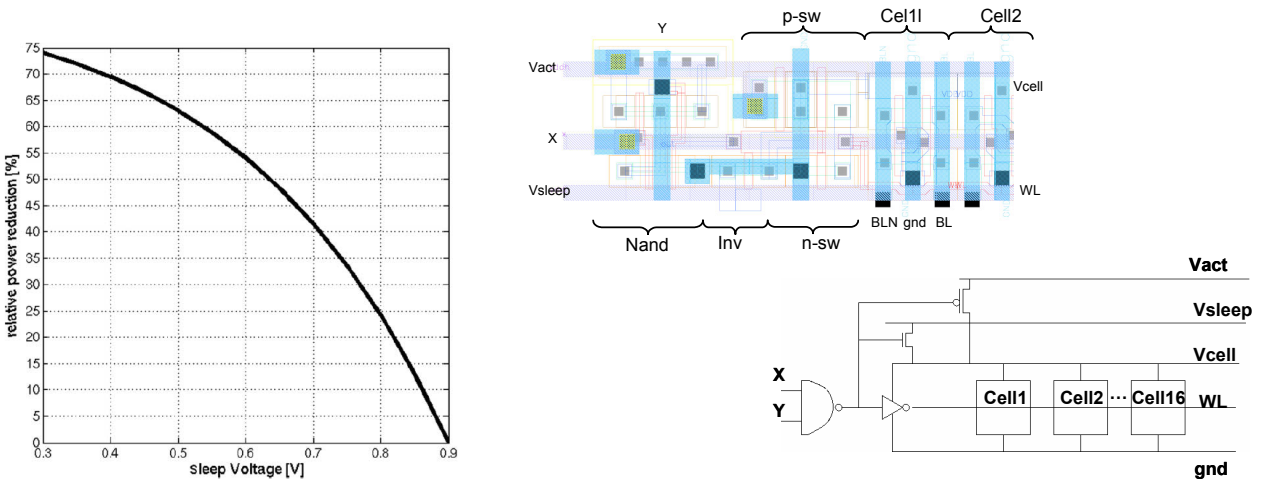


Figure 32.1.5: Fine grain leakage controlled SRAM: (a) theoretical power reduction in the matrix as a function of the lower sleep voltage(8Kb matrix, 16b words); (b) layout and schematic of one word in the matrix.